Department of Economics

Working Paper Series

Adverse Selection and Moral Hazard: Quantitative Implications for Unemployment Insurance

12-004 | David L. Fuller
Concordia University and CIREQ

Department of Economics, 1455 De Maisonneuve Blvd. West, Montreal, Quebec, Canada H3G 1M8
Tel 514–848–2424 # 3900 · Fax 514–848–4536 · econ@alcor.concordia.ca · alcor.concordia.ca/~econ/repec

Adverse Selection and Moral Hazard: Quantitative Implications for Unemployment Insurance

David L. Fuller*

Department of Economics

Concordia University

October 17, 2013

Abstract

A model of optimal unemployment insurance with adverse selection and moral hazard is constructed. The model generates both qualitative and quantitative implications for the optimal provision of unemployment insurance. Qualitatively, for some agents, incentives in the optimal contract imply consumption increases over the duration of non-employment. Calibrating the model to a stylized version of the U.S. economy quantitatively illustrates these theoretical predictions. The optimal contract achieves a welfare gain of 1.94% relative to the current U.S. system, an additional 0.87% of gains relative to a planner who ignores adverse selection and focuses only on moral hazard.

Keywords: unemployment insurance, non-participation, adverse selection, moral hazard, dynamic contracts

JEL classification: C61, D82, E61, J64, J65

*1455 de Maisonneuve Boulevard West, Montreal, Quebec H3G 1M8, Canada, (514)848-2424 extension 3919; david.fuller@concordia.ca

I would like to thank thank Paul Gomme, Matt Mitchell, B. Ravikumar, Chris Sleet, Gustavo Ventura, Galina Vereshchagina, and Yuzhe Zhang, as well as many conference and seminar participants for helpful comments and encouragement. I would also like to thank the Associate Editor Yongsung Chang and an anonymous referee for comments that greatly improved the paper.

1 Introduction

Since the work of Shavell and Weiss (1979) and Hopenhayn and Nicolini (1997) a large literature examining the optimal provision of unemployment insurance (UI) has developed. This literature has focused on circumventing the moral hazard problem, which arises when agents exert unobservable effort to find or retain a job. Given the emphasis on moral hazard, existing models of optimal unemployment insurance focus on optimizing the transitions between unemployment (defined as actively searching for a job) and employment. While the literature contains many interesting results, two are robust: an agent's consumption should decrease over the duration of an unemployment spell (Kocherlakota (2004) and Hagedorn et al. (2010) are two exceptions), and there exist relatively large welfare gains from adopting the optimal contract.

Circumventing moral hazard requires the UI agency to know the utility costs of exerting effort. The analysis in this paper considers the case where agents are heterogeneous in these costs, which remain private information for the agent; as a result, there exists moral hazard *and* adverse selection.[1] The contribution of this paper to the optimal UI literature is twofold. First, analytically it is shown that the inclusion of adverse selection matters for the dynamics of UI benefits. Quantitatively, a plausible, stylized model of the U.S. system verifies that the dynamics of benefits are different from the pure moral hazard case. Second, the stylized quantitative analysis also shows that there exist relatively large additional welfare gains from considering both adverse selection and moral hazard.

Specifically, the adverse selection in my model arises from unobservable, idiosyncratic preference shocks, affecting the marginal rate of substitution between consumption and leisure. As in Wang

[1] Alternatively, one can substitute the nomenclature "hidden information" in place of "adverse selection", and "hidden action" in place of "moral hazard".

and Williamson (2002), agents move in and out of employment, exerting unobservable search effort while unemployed, and job retention effort while employed. In some periods, agents may receive a taste shock that makes exerting effort too costly; consequently the agent might prefer not to search in that period.

Large transitions from unemployment (searching) to non-participation (not-searching) in U.S. data motivate the inclusion of heterogeneity in the utility cost of effort (and thus adverse selection). For example, for the period 1994-2004, on average, 30% of unemployed individuals moved to non-participation.[2] The standard moral hazard model of UI, however, does not capture this transition (Andolfatto and Gomme (1996), Garibaldi and Wasmer (2005) Pries and Rogerson (2009), Gautier et al. (2009), and Engelhardt and Fuller (2012), among others, do include non-participation, but do not analyze the optimal UI scheme). Qualitatively, the model captures this transition, and using the data allows a quantitative illustration of the key implications of including adverse selection. While this quantitative analysis does not account for the full heterogeneity of the group of non-participants, it does suggest important implications from including this dimension in the analysis of optimal UI.

First, incorporating adverse selection has implications for the dynamics of benefits, which remain relevant for policy. Consumption in the optimal contract does not necessarily decrease over the duration of non-employment, as pure moral hazard models suggest. Analytically, there exists a positive probability of consumption increasing over the duration of non-employment, and the quantitative analysis shows that this result is relevant for policy.[3] Specifically, the optimal contract implies that after an initial spell of unemployment, an agent who transitions to non-participation and remains there (under the U.S. system), should receive increasing consumption during the spell

[2]The 30% is a quarterly average of monthly transition probabilities. See Section 2 for details of my calculations.
[3]All proofs and supplementary appendices are available in the Online Appendix, or at Science Direct.

of non-participation.

Increasing consumption occurs because to efficiently allocate effort, the UI authority asks each agent to report his desire to remain in the labor force (i.e. their idiosyncratic taste shock). Truth-telling incentives imply an agent with more attachment receives relatively more consumption today in exchange for less consumption in all future periods, while an agent reporting a low attachment receives less consumption today, but is promised relatively more in the future. Although the planner may want to provide incentives to exert effort, by reducing future consumption when non-employed, near the participation constraint, there exists little room for punishment and truth-telling incentives dominate.

To quantitatively illustrate the main results, the model is calibrated to U.S. data and a stylized version of the U.S. system of unemployment insurance, focusing on capturing the unemployment to non-participation transition. In the optimal contract the planner efficiently allocates effort, which is a discrete choice: exert effort or not. In the baseline calibration, the transitions from unemployment to non-participation occurring under the U.S. system are not efficient; i.e. the optimal contract recommends effort from the same agents.

To further evaluate the efficacy of the current U.S. system, imagine switching from the U.S. system to the optimal contract. This switch yields a welfare gain of almost 2% in consumption equivalent terms. To determine what role accounting for the non-participation dimension has in these welfare gains, the optimal contract is compared to an allocation from a "naïve" planner, who ignores adverse selection and focuses only on moral hazard. For the baseline calibration, the optimal contract provides an additional 0.87% in consumption equivalent welfare gains. The comparison is also made over sub-sets of the population, based on age. There exist even larger additional gains from the optimal contract for the youngest cohort (16-24), 1.87%, while much smaller additional gains are available for the oldest cohort (55 and older), 0.30%.

As noted above, the lower bound on expected utility promises plays a role in the increasing consumption result. Both Atkeson and Lucas (1995) and Pavoni (2007) impose such a lower bound in models of optimal unemployment insurance, and analyze how it circumvents the well-known immiseration result. In the quantitative analysis, the lower bound is "endogenized" by studying a general equilibrium version of the planning problem, similar to the analysis in Atkeson and Lucas (1995) and Wang and Williamson (2002). The general equilibrium version determines the steady state distribution over lifetime promised utility induced by the optimal contract. The value of the lower bound is linked to the autarky value of an agent in the U.S. economy, and is further endogenized it by examining the range of autarky values that admit a balanced budget equilibrium in the planning problem.

Temporary shocks remain necessary to generate the decrease in search intensity behind the non-participation transition. In addition to temporary shocks, i.i.d. or more persistent, there may also exist permanent differences between unemployed agents; for example, agents may have innate differences in search "ability." Hagedorn et al. (2010) analyze the optimal unemployment insurance scheme in such an environment, extending the models of Shavell and Weiss (1979) and Hopenhayn and Nicolini (1997), and similar to this paper, have an environment with both adverse selection and moral hazard. Since these permanent differences likely exist, Hagedorn et al. (2010) represents a careful analysis of their consequences for optimal unemployment insurance schemes. The important distinction between the model in Hagedorn et al. (2010) and this paper is that the former does not capture the non-participation dimension. Moreover, Hagedorn et al. (2010) do not offer a comparison of the optimal scheme to the current U.S. policy, which is made here.

The analysis in this paper also relates to the dynamic adverse selection literature.[4] Atkeson and Lucas (1992) represents the paper from this literature most closely related to mine, as they also analyze unobservable, i.i.d., taste shocks. As in this paper, the planner in Atkeson and Lucas (1992) offers high consumption today in exchange for lower future consumption, when reporting a high marginal utility of consumption. A similar provision of incentives also occurs in Thomas and Worral (1990), who analyze dynamic adverse selection in a model with i.i.d., unobservable idiosyncratic income shocks. Farhi and Werning (2007) analyze a model similar to Atkeson and Lucas (1992), showing how incorporating generational concerns circumvents immiseration. Although the focus remains on distributional issues, the management of incentives across taste shocks in Farhi and Werning (2007) is similar to that found in other dynamic adverse selection environments, including in this paper.

Finally, the analysis of the optimal contract in this paper follows the descendants of Hopenhayn and Nicolini (1997) by assuming the planner controls an agent's consumption. Thus, unemployment benefits and consumption are identical. A recent strand of the literature on optimal unemployment insurance removes the planner's control over an agent's consumption (For example, see Werning (2002), Kocherlakota (2004), Shimer and Werning (2006, 2008) , and Mitchell and Zhang (2010)). Specifically, these papers allow agents to borrow and save, and in some cases this activity takes place behind the planner's back.

While there exist interesting effects from the presence of hidden savings, declining consumption over a spell of unemployment still obtains in these models (Kocherlakota is the only exception, who finds constant consumption). In certain cases, *benefits* may increase over the duration of an unemployment spell (for example see Shimer and Werning (2006) or Mitchell and Zhang (2010)),

[4]Since this literature is large, only those papers closely related are discussed.

but since consumption and benefits are separated when agents can save and borrow, increasing benefits does not imply increasing consumption. Indeed, in the cases where benefits increase over the duration of unemployment, a decreasing consumption profile still obtains.

2 Evidence on Search Intensity

To produce the transition from unemployment to non-participation, taste shocks affecting the marginal rate of substitution between consumption and leisure are incorporated. Differences in search "ability," or discouraged job searchers represent two other plausible explanations for this observed transition. Indeed, one strand of the literature on unemployment insurance has studied the effects of human capital depreciation on the optimal unemployment insurance scheme (for examples, see Pavoni (2009) and Shimer and Werning (2006)). This section provides evidence on the search intensity of unemployed agents in the U.S., motivating the modeling choices in this paper.

First consider the group of agents who have transitioned from unemployment to non-participation. To perform this analysis (and to calibrate several parameters of the model later), the same method as Shimer (2012), who uses CPS (Current Population Survey) data to calculate the flows in and out of the different labor market states, is used. From this data, on average over the period $1994-2004$, around 30% of agents who were unemployed and searching last period, decide to stop searching in the current period, and thus transition to non-participation.[5] These numbers reflect quarterly averages of monthly transition probabilities.

Table 1 presents data on the reasons agents gave for the decision to stop searching. The data

[5] According to the CPS classification, "searchers" are defined as unemployed workers, or those who have actively searched for work in the previous four weeks. "Non-searchers" are those considered not-in-the-labor-force (i.e. non-participants), who have not actively searched in the previous four weeks.

represent yearly averages from the monthly CPS in 2005. The group in this table consists of persons who searched for work in the past year, are currently not searching, but want a job and are available for work. A large majority of this group (72%) has stopped searching for reasons other than discouragement, supporting the modeling choice used in this paper.

3 Environment

This section describes succinctly the economic environment defining the dynamic contracting problem. It describes preferences, the model timing, and the agent's effort decision.

3.1 Preferences

There exists a risk-neutral planner offering an unemployment insurance contract to a risk-averse agent with preferences given by:

$$E_0 \sum_{t=0}^{\infty} \beta^t [\theta_t \mathrm{v}(c_t) - \nu(a_t)] \tag{1}$$

Here c_t represents consumption, $\mathrm{v}(c_t)$ the per-period utility from consumption, $\nu(\cdot)$ measures the utility cost from exerting effort a_t, and θ_t represents a preference shock. Assume that consumption remains observable, while θ_t and a_t represent private information for the agent. Effort is a discrete decision, $a_t \in \{0, 1\}$; either exert effort, $a_t = 1$, or not, $a_t = 0$. The preference shock θ_t is i.i.d., with $\theta \in \{\theta_l, \theta_h\}$, and $\theta_h > \theta_l > 0$. The probability of receiving the shock θ_h is given by μ. Also assume the per period utility function, $\mathrm{v}(\cdot)$, is increasing, C^2, and strictly concave with $\mathrm{v}'(0) = \infty$. The utility cost of effort function, $\nu(\cdot)$, is $\nu(1) = \nu$, and $\nu(0) = 0$.

It is convenient to think of the planner as offering agents a current level of utility, v_t, instead of consumption. That is, $\mathrm{v}(c_t) = \mathrm{v}_t$. Given the assumption of strict concavity of the utility function, define the inverse utility function $C(\cdot)$, which is strictly convex and increasing.

3.2 Timing

An agent enters each period with his employment status from last period, j, and his expected lifetime utility promise, w, known. If the agent was employed last period, $j = e$, while $j = u$ represents the unemployed (i.e. non-employed) state. During the period, the agent chooses a reporting strategy (report either θ_l or θ_h) and makes an effort decision. Based on this effort, there is a probability $\pi_j(a_t)$ the agent is employed in the current period ($1 - \pi_j(a_t)$ he is unemployed), where $\pi_j(1) = \pi_j$ and $\pi_j(0) = 0$. If employed, the agent produces $y > 0$ units of output, which the planner confiscates, and unemployed agents produce nothing. The contract specifies utility for the current period, the agent's future lifetime utility promise, depending on the reported θ and whether the agent finds employment during the period or not. The planner offers the following contract to an agent entering the period in employment state j:

$$\tau_j = \left\{ a_j^k, \mathrm{v}_j^k(u), \mathrm{v}_j^k(e), w_j^k(u), w_j^k(e), \right\}, k = h, l \tag{2}$$

For example, consider an agent who enters the period employed ($j = e$), and reports θ_h. For this agent, a_e^h is the recommended level of effort, and $\mathrm{v}_e^h(u)$ represents the current period utility if unemployed at the end of the period, while $\mathrm{v}_e^h(e)$ is the current period utility allocated to this agent if he remains employed at the end of the period. Similarly, $w_e^h(u)$ represents the future lifetime utility promise if this agent is unemployed at the end of the period, and $w_e^h(e)$ the future lifetime utility promise if he remains employed.[6] The timing used here follows closely that used in Wang

[6]The allocations also depend on w and should be denoted $\mathrm{v}_j^k(e, w)$ and $w_j^k(e, w)$, for agents leaving employed, and similarly for those leaving unemployed. While this dependence is omitted in the notation, it is implied everywhere.

and Williamson (2002), who also model both the unemployment-to-employment and employment-to-unemployment transition.

3.3 Agent Decision

Before presenting the planning problem, it is useful to describe an agent's decision problem. Denote the expected lifetime utility from the contract τ_j, for an agent receiving the shock θ_i, reporting the shock θ_k, and exerting effort a, by $V_j(\theta_i, \theta_k, a)$. This is given by,

$$V_j(\theta_i, \theta_k, 1) = \theta_i[\pi_j \mathrm{v}_j^k(e) + (1-\pi_j)\mathrm{v}_j^k(u)] - \nu + \beta \left[\pi_j w_j^k(e) + (1-\pi_j) w_j^k(u)\right] \tag{3}$$

$$V_j(\theta_i, \theta_k, 0) = \theta_i \mathrm{v}_j^k(u) + \beta w_j^k(u) \tag{4}$$

This agent exerts effort if $V_j(\theta_i, \theta_k, 1) \geq V_j(\theta_i, \theta_k, 0)$, which occurs when,

$$\theta_i \pi_j \left[\mathrm{v}_j^k(e) - \mathrm{v}_j^k(u)\right] + \beta \pi_j \left[w_j^k(e) - w_j^k(u)\right] \geq \nu \tag{5}$$

Equation (5) illustrates how the taste shock affects an agent's effort decision. Under the timing assumption in this model, effort exerted in the current period affects the current period employment status. This feature is important for recognizing how the taste shock affects effort. Exerting effort implies a better chance for higher consumption from employment. When receiving the θ_h shock, this higher consumption is appealing, since the marginal utility from consumption is high. On the other hand, an agent receiving the θ_l shock may decide not to exert effort as the marginal benefit remains low relative to the utility cost of exerting effort.

Allowing the taste shock to directly hit the utility cost of effort represents an alternative (and perhaps more natural) way to model the effort decision. In this case, preferences are given by

$\mathrm{v}(c) - \theta_t \nu(a)$. This specification allows the timing of the employment shock to only affect next periods' state, similar to Hopenhayn and Nicolini (1997). While more natural, with a discrete choice of effort technical issues arise, making the analytical results less clean. Briefly, with a discrete choice of effort (and only two effort values) the planner may not be interested in learning the value of the taste shock; that is, a "pooling" contract remains possible. This can occur when the planner prefers both reports to exert effort. In this case, characterizing the main analytical result (Proposition 1 below) is difficult. The details of this case are discussed further in the Online Appendix, where the equivalent of Proposition 1 is proved for the case where the planner does not request effort from one report.

Moreover, Section 6.4 shows that the welfare gains are similar under either specification of preferences, suggesting that quantitatively the main results are not affected by this particular specification. Since the qualitative and quantitative results are unaffected, the analysis proceeds with the preferences analyzed in this section where the theoretical results are sharper.

4 Planning Problem

The social planner attempts to find the lowest cost, incentive compatible contract. The expected cost to the planner of providing the contract τ_j, to an agent with employment status j, is denoted by $G_j(w)$. The generalized problem is presented, with lotteries over the different combinations of effort, which provides convex cost functions.[7]

For a given j, there exist four permutations of effort combinations across values of θ. Denote (a^l, a^h) the effort choices for an agent in state j. Depending on the choice of effort, the planner

[7]Since effort, remains a discrete choice variable for the planner, in general the problem may not be convex; as a result, lotteries over choice variables may be necessary. The analysis of the generalized problem employs similar techniques to that in Hopenhayn and Nicolini (2009), who study a pure moral hazard problem with discrete effort choices.

faces four possible value functions. Let $\vec{a}_i = (a_i^l, a_i^h)$ denote the planner's allocation of effort for $i \in \{1,2,3,4\}$, where $\vec{a}_1 = (0,1), \vec{a}_2 = (1,1), \vec{a}_3 = (0,0)$, and $\vec{a}_4 = (1,0)$. The analysis below only considers $i \in \{1,2,3\}$, since $\vec{a}_4$ is never optimal. Further let $G_j^i(w)$ denote the value function for the effort allocation $\vec{a}_i$. Given the choice of effort, $\vec{a}_i$, the planner chooses an allocation τ_j^i for an agent in employment state j. The allocation specifies

$$\tau_j^i = \left\{ \mathrm{v}_j^h(u,i), \mathrm{v}_j^h(e,i), \mathrm{v}_j^l(u,i), \mathrm{v}_j^l(e,i), w_j^h(u,i), w_j^h(e,i), w_j^l(u,i), w_j^l(e,i) \right\} \tag{6}$$

Finally, let A_k^i denote the set of feasible effort choices for an agent reporting the shock θ_k under the planner's effort allocation $\vec{a}_i$. For example, under $\vec{a}_1$, $A_h^1 = \{0,1\}$ and $A_l^1 = \{0\}$. Thus, if an agent reports θ_l, even if he received the shock θ_h, $\tilde{a} = 0$ is the only feasible choice of effort. Since this agent cannot report employment to the planner, exerting effort is never optimal.

For a given choice of effort, $\vec{a}_i$, denote $\pi_j(a_i^k)$ by $\pi_j^{k,i}, k = h, l$. Further let E_θ denote the expectation over the two values of θ_k. Then, the planner solves:[8]

$$G_j^i(w) = \min_{\tau_j^i} E_\theta \left\{ \pi_j^{k,i} \left[C(\mathrm{v}_j^k(e)) - y + \beta G_e(w_j^k(e)) \right] + (1 - \pi_j^{k,i}) \left[C(\mathrm{v}_j^k(u)) + \beta G_u(w_j^k(u)) \right] \right\} \tag{7}$$

$$\text{s.t.} \quad w = \mu V_j(\theta_h, \theta_h, a_i^h) + (1-\mu) V_j(\theta_l, \theta_l, a_i^l) \tag{8}$$

$$V_j(\theta_h, \theta_h, a_i^h) \geq V_j(\theta_h, \theta_l, \tilde{a}), \forall \tilde{a} \in A_l^i \tag{9}$$

$$V_j(\theta_h, \theta_h, a_i^h) \geq V_j(\theta_h, \theta_h, \tilde{a}), \forall \tilde{a} \in A_h^i \tag{10}$$

$$V_j(\theta_l, \theta_l, a_i^l) \geq V_j(\theta_l, \theta_h, \tilde{a}), \forall \tilde{a} \in A_h^i \tag{11}$$

$$V_j(\theta_l, \theta_l, a_i^l) \geq V_j(\theta_l, \theta_l, \tilde{a}), \forall \tilde{a} \in A_l^i \tag{12}$$

[8] The dependence of allocations on i, the particular effort choice, is suppressed. For example, $\mathrm{v}_j^k(e,i)$ is written as $\mathrm{v}_j^k(e)$.

Equation (8) represents the promise keeping constraint. The incentive compatibility constraints, ensuring an agent receiving the shock θ_h truthfully reports are given by equation (9). These truth-telling constraints illustrate the interaction of the adverse selection and moral hazard frictions, as an agent can simultaneously lie and shirk. Equation (10) requires the agent truthfully reporting θ_h to prefer exerting effort a_i^h. Equations (11)-(12) represent the equivalent incentive constraints for an agent receiving the shock θ_l.

Finally, following Atkeson and Lucas (1995), Wang and Williamson (2002), and Pavoni (2007), a lower bound on future expected utility promises is imposed,

$$w_j^k(i) \geq \underline{w}, \quad j = e, u; k = h, l; i = e, u \tag{13}$$

This lower bound can be interpreted as a participation constraint: the planner must deliver at least some minimum level of consumption to an agent, or he can always refuse the contract. That is, the lower bound represents the value of autarky for the agent. The quantitative analysis in Section 5 provides additional links between current U.S. unemployment insurance system, the lower bound, and the main theoretical results of the paper presented below.

To convexify the problem, the planner chooses lotteries over the different choices of $\vec{a}$. Denote the choice of lottery by $\vec{q} = \{q_i\}_{i=1}^3$, where q_i is the probability of effort choice $\vec{a}_i$. The Online Appendix provides details on the more general problem and a formal proof of convexity.

4.1 Properties of the Optimal Contract

The main analytical result of the paper, that with positive probability, an agent's consumption *increases* over the duration of an unemployment (non-employment) spell is now presented.

Proposition 1 *There exists a $\delta > 0$ such that an agent starting an unemployment spell with promised lifetime utility $w_0 \in [\underline{w}, \underline{w} + \delta)$, receives increasing consumption over the length of the unemployment spell with positive probability.*

While counterintuitive at first, the intuition behind this result is actually straightforward. To understand this, consider the planner's allocation to an agent leaving the period unemployed. Efficient provision of incentives implies that the planner allocates more current utility to an agent reporting θ_h; i.e. $\mathrm{v}_j^h(u) > \mathrm{v}_j^l(u)$. Then, to induce truth-telling from an agent receiving the shock θ_l, the planner sets $w_j^l(u) > w_j^h(u) \geq \underline{w}$. Now suppose an agent starts an unemployment spell with initial promised utility w_0 near $\underline{w}$. Further suppose this agent remains unemployed for T periods, and receives a sequence of low shocks, $\{\theta_l, \theta_l, ..., \theta_l\}$ (this sequence of shocks occurs with positive probability). Since $w_u^l(u) > w$ for an interval, the agent starting near the lower bound experiences increasing promised utility during the unemployment spell; as a result, consumption also increases.

4.2 General Equilibrium Planning Problem

Thus far, the analysis of the planner's problem is "partial equilibrium": the planner meets only one agent and provides him with an allocation; there is no balanced budget constraint. This section introduces the general equilibrium version of the planning problem, following Atkeson and Lucas (1995) and Wang and Williamson (2002).

In the general equilibrium version, there exists a planner for each level of expected utility. These planners trade with each other at intertemporal prices, $\{P_t\}_{t=0}^{\infty}$, given by $P_t = \prod_{s=0}^{t-1} p_s$, where $p_t \in (0, 1)$. Only steady states are examined, so that $p_t = p, \forall t$. This p replaces β in the planner's objective function in equation (7).

Denoting by $\psi_{t,j}$ the current distribution of agents (in employment state j) across expected utility states in period t, notice how the planner's choice of $w_j^k(i)$ induces a new distribution across expected utility states next period, $\psi_{t+l,j}(w)$. That is, there exists an operator Γ, mapping distributions of expected utility into distributions of expected utility, $\Gamma\psi_{t,j}(w) = \psi_{t+1,j}(w)$. Again, the analysis here focuses on steady states where $\Gamma\psi(w) = \psi(w)$, with $\psi(w)$ representing the vector of unique invariant distributions. Given intertemporal prices, p, the following aggregate resource constraint must be satisfied:

$$\sum_j \int \left\{ E_\theta \left[\pi_j^k (C(\mathrm{v}_j^h(e)) - y) + (1 - \pi_j^k) C(\mathrm{v}_j^h(u)) \right] \right\} d\psi_j(w) = 0 \tag{14}$$

where E_θ is the expectation over the two values of θ_k and π_j^k is the probability of employment for an agent currently in employment state j reporting the shock θ_k, given the planner's allocation of effort for that report. The resource constraint then, specifies that total consumption must equal total output.

This general equilibrium planning problem provides a quantitative foundation for what the lower bound represents, and how it relates to the parameters of the existing unemployment (and more general social) insurance scheme(s) in the U.S. The quantitative analysis below links the lower bound to the autarky value of an agent in the U.S. system, and the range of autarky values that admit a balanced budget equilibrium in the planning problem are considered.[9] Thus, the maximum and minimum "punishments" that permit a balanced budget in the optimal allocation are determined.

[9] Since $p \in (\beta, 1)$, for a certain $\underline{w}$ there may not exist an equilibrium in the planning problem; that is, for small enough $\underline{w}$ aggregate output always exceeds aggregate consumption, and for large enough $\underline{w}$ aggregate consumption always exceeds output, regardless of p.

5 Quantitative Analysis

Proposition 1 shows that the dynamics of benefits may be quite different from a standard moral hazard model. Quantitatively illustrating the implications of adverse selection represents the primary goal of this section. The observed transition from unemployment to non-participation (see Section 2) motivates the inclusion of heterogeneity and helps quantify the extent of the underlying the adverse selection problem. Attention is further restricted to a subset of the non-participant population. Although the exercise does not account for the full heterogeneity of the non-participant population, there exist large additional welfare gains from considering adverse selection in addition to moral hazard.

To calibrate the model, the relevant parameters are chosen so the implied behavior of agents facing a stylized version of the current U.S. unemployment insurance scheme matches U.S. data. Then, given these parameters, the optimal contract is computed, the characteristics of consumption over the duration of non-employment are illustrated, and the welfare gains from accounting for adverse selection are quantified.

5.1 U.S. System

Consider the problem of an agent faced with the current U.S. unemployment insurance scheme. This scheme varies across states, but all represent some variation of a constant benefit for a fixed length of time, typically 26 weeks. The time period is one quarter; thus, in the stylized model presented below, agents receive benefits for two periods, which translates into two quarters, or 26 weeks. For the calibration period, 1994-2004, total UI expenditures in the U.S. averaged \$37.9 billion per year, in 2011 dollars (source: U.S. Department of Labor).

As in the planning problem, if an agent in employment state j last period exerts effort, he is employed today with probability π_j. If employed in the current period, the agent consumes his output y, while if non-employed he collects benefits, which depend on the length of non-employment. If the agent has been non-employed for less than two periods, he receives consumption of $b = 0.66y$. While higher than some estimates for the replacement rate, this appears reasonable given the lack of precautionary savings in the model.[10] The unemployment insurance system is subject to a balanced budget constraint, with lump-sum taxes, denoted by τ_b, levied on employed agents. Once these benefits have expired, the agent collects some minimum level of consumption, d. This level is calibrated below, based on the autarky value of an agent in the U.S. system and the existence of an equilibrium in the planning problem.

An agent can be in one of five states: employed or one of four non-employment states. To reduce notation, only two separate non-employment states are presented: $j = u$ and $j = nb$, which are differentiated based on benefit eligibility. Depending on search effort, however, there exist four potential non-employment states. More specifically, an agent can be in one of the following states: Employed: $j = e$; Unemployed (searching), receiving benefits: $j = u$; Unemployed (searching), no benefits ($j = nb$); Non-participant (not searching), receiving benefits: $j = u$; Non-participant (not searching), no benefits: $j = nb$.

Denote the expected lifetime utility of being in state j entering the period, receiving the shock θ_k, and exerting effort a, as $V_j^{US}(\theta_k, a)$. Furthermore, let $V_j^{US}(\theta_k) = \max\left\{V_j^{US}(\theta_k, 1), V_j^{US}(\theta_k, 0)\right\}$.

[10] Replacement rate estimates vary across the UI literature. Meyer (1990) estimates $b = 0.66$ for a sample of 12 states. Hopenhayn and Nicolini (1997) use 0.66, while Wang and Williamson (2002) use the value $b = 0.5y$, which is taken from Davidson and Woodbury (1998). Since employed agents in my model do not incur any utility cost of working, one could argue for a replacement rate higher than $b = 0.66$. That is, given the unemployed suffer a utility loss from search, unemployed utility may be too far below employed utility with $b = 0.66$. Hall and Milgrom (2008) make a similar argument when they set unemployed consumption equal to 85% of employed consumption.

An agent entering the period in employment state $j \in \{e, u, nb\}$ faces the following problem.

$$V_j^{US}(\theta_k, 1) = \theta_k \left[\pi_j \mathrm{v}(y - \tau_b) + (1 - \pi_j)\mathrm{v}(b_j)\right] - \nu + \beta E_{\theta'} \left\{\left[\pi_j V_e^{US}(\theta') + (1 - \pi_j) V_i^{US}(\theta')\right]\right\} \quad (15)$$

$$V_j^{US}(\theta_k, 0) = \theta_k \mathrm{v}(b_j) + \beta \left[\mu V_i^{US}(\theta_h) + (1 - \mu) V_i^{US}(\theta_l)\right] \quad (16)$$

Here $i \in \{u, nb\}$ denotes the possible non-employment states for the agent currently in state j, b_j the corresponding unemployment benefits, with $\pi_{nb} = \pi_u$. If an employed agent loses the job, $i = u$, and $b_e = b$. For an agent in state $j = u$, if he remains unemployed at the end of the period $i = nb$ and $b_u = b$. An agent in state $j = nb$ remains in this state if unemployed at the end of the period, and having exhausted unemployment benefits, $b_{nb} = d$.

5.2 Autarky

Now consider the agent's problem in autarky. Under autarky, the number of non-employment states is reduced from four to two; an agent is either non-employed and searching, or non-employed and not searching. Autarky is defined as never having access to the unemployment benefits, b, but only the minimum level of consumption d. Thus, autarky is distinct from $V_{nb}^{US}(\theta_k)$ above. Denoting $V_j^A(\theta_k) = \max\left\{V_j^A(\theta_k, 1), V_j^A(\theta_k, 0)\right\}$ the value functions under autarky, the system is defined by:

$$V_j^A(\theta_k, 1) = \theta_k \left[\pi_j \mathrm{v}(y) + (1 - \pi_j)\mathrm{v}(d)\right] - \nu + \beta E_{\theta'} \left\{\left[\pi_j V_e^A(\theta') + (1 - \pi_j) V_{nb}^A(\theta')\right]\right\} \quad (17)$$

$$V_j^A(\theta_k, 0) = \theta_k \mathrm{v}(d) + \beta \left[\mu V_{nb}^A(\theta_h) + (1 - \mu) V_{nb}^A(\theta_l)\right] \quad (18)$$

To calibrate the lower bound, $\underline{w}$, consider the aforementioned features of the general equilibrium version of the planner's problem. Specifically, the lower bound is set to the expected utility an agent receives in autarky, which is given by:

$$\underline{w} = \mu V_{nb}^{A}(\theta_h) + (1-\mu)V_{nb}^{A}(\theta_l) \tag{19}$$

Depending on the lower bound, an equilibrium in the planning problem may or may not exist. The value of d is taken as those which admit a balanced budget equilibrium, and the results over this range are presented. For the baseline case, $d = 0.25y$. Given this and the value of b, the parameters imply that over a 60 month spell of non-employment, the average replacement rate is 0.29. This approximately matches data reported by the OECD on replacement rates over this length of non-employment spell (OECD data implies an average replacement rate of 0.36).[11] Values for d between 0.25 and 0.37 provide the range of $\underline{w}$ that admit an equilibrium in the planning problem.

The quantitative analysis below makes a comparison of the optimal contract with the U.S. system for currently employed agents. The U.S. system described above delivers expected lifetime utility to an agent in employment state $j = e$ of

$$w_{US} = \mu V_{e}^{US}(\theta_h) + (1-\mu)V_{e}^{US}(\theta_l) \tag{20}$$

5.3 Parameters

The model described above leaves the parameters $\beta, \nu, \pi_u, \pi_e, \mu, \theta_h$, θ_l, the utility function $\mathrm{v}(c)$, and the equilibrium tax rate, τ_b, to be specified. The per-period utility function is $\mathrm{v}(c) = \frac{c^{1-\sigma}-1}{1-\sigma}$, with $\sigma = 1$, or $\mathrm{v}(c) = \log(c)$, consistent with Wang and Williamson (2002). Further, several of the parameters above can be recovered directly from U.S. data. For β, assume $\beta = \frac{1}{1+r}$, where r represents the risk free interest rate. Since the time period is one quarter, $r = 0.04$ (per annum)

[11] http://www.oecd.org/els/benefitsandwagesstatistics.htm

implies $\beta = 0.99$. The value of π_j is obtained using data on the flows in and out of unemployment. For the period from 1994-2004, the observed flows in and out of unemployment imply $\pi_u = 0.330$ and $\pi_e = 0.982$, again using CPS data along with the method cited in Section 2.

Given the parameters listed above, the remaining parameters, ν and the distribution of θ, should be chosen such that the model generates the unemployment to non-participation transition observed in the data. Doing so requires agents receiving the θ_l shock to prefer not searching while non-employed. Under this interpretation, the probability of receiving θ_l, $1 - \mu$, represents the probability an agent transitions from unemployment to non-participation. Thus, ν,θ_h, θ_l, and μ are chosen such that the agents receiving θ_h always decide to search, and those receiving the shock θ_l, decide not to search (while non-employed).

The observed flows of unemployed persons to non-participation determine $1 - \mu$. Letting p_{ns} denote the probability an unemployed agent exits to the non-participation state, on average for the period from 1994-2004, $p_{ns} = 0.295$. To rule out agents who have received more persistent or permanent shocks, only agents who move from unemployment to non-participation but still want a job and are available for work are considered. This represents a group of agents who may have received a temporary shock, and are still interested in actively participating in the labor market. Specifically, for the period 1994-2004, on average 7.2% percent of those considered non-participants wanted a job now. Combined with $p_{ns} = 0.295$, this implies $1 - \mu = 0.021$.

Choosing ν, θ_h, and θ_l represents the final step in the calibration. Towards this end, the value of the shocks are normalize: $E(\theta) = 1$. Then, choose the *smallest* ratio, $\frac{\theta_h}{\theta_l}$, such that the agent receiving the shock θ_h always prefers searching while non-employed, and the agent receiving θ_l always prefers not searching. In terms of the agents' decision problem when facing the U.S. scheme, this can be accomplished if: (i) $V_u^{US}(\theta_h, 1) = V_u^{US}(\theta_h, 0)$, and (ii) $V_{nb}^{US}(\theta_l, 1) + \varepsilon = V_{nb}^{US}(\theta_l, 0)$, for

$\varepsilon > 0$ small, both hold.

Given the aforementioned requirements for θ, the ratio $\frac{\theta_h}{\theta_l}$ can be solved for directly. This ratio is given by (see the Online Appendix for details):

$$\frac{\theta_h}{\theta_l} = \frac{[\mathrm{v}(y - \tau_b) - \mathrm{v}(d)]}{[\mathrm{v}(y - \tau_b) - \mathrm{v}(b)]} \tag{21}$$

This uniquely determines the smallest ratio $\frac{\theta_h}{\theta_l}$ such that the θ_h agent always searches and the θ_l agent never does (while non-employed). Given the normalization, $E(\theta) = 1$, this ratio pins down both θ_h and θ_l. Finally, given these values for θ_h and θ_l, ν is set to ensure the θ_h agent remains indifferent between exerting effort or not. In the baseline parametrization, $\dfrac{\theta_h}{\theta_l} = 3.49$ and $\nu = 0.98$.

With the parameters specified, the analysis now turns towards examining the qualitative and quantitative implications of the optimal contract, when compared to the current U.S. system.

5.4 Optimal Allocation of Effort

In the baseline parametrization, the optimal effort allocation at $w = w_{US}$ is $\vec{q} = (0, 1, 0)$; all non-employed agents exert effort. All employed agents also exert effort in the optimal allocation. The implication of this allocation of effort is that the transition from unemployment to non-participation occurring under the current U.S. system is not efficient; the planner prefers to have those non-participants under the U.S. system (agents receiving the shock θ_l) to continue searching.

Table 2 compares the changes in transitions rates between employment states from the current U.S. system, naïve allocation (described below), and optimal contract, respectively. Since the optimal contract prefers all reports to exert effort, the transition rate from unemployment to employment (U-E) increases, and the transition from unemployment to non-participation (U-O) decreases. Of course, in reality, the parameters π_u and π_e are also likely to change under the optimal contract, if taking into account the intensive margin of effort. Since the model only includes the extensive margin, however, these parameters are fixed; as a result, these changes in transition rates understate the likely impact of the optimal contract.

5.5 Increasing Consumption and Non-participation

To illustrate the relevance of Proposition 1 for current policy, suppose an agent starts a spell of non-employment with $w = w_0$, and remains non-employed for T periods. Consider an agent who begins the non-employment spell with the mean level of expected utility for employed agents in equilibrium under the optimal allocation; denote this level by $\bar{w}_E$. The agent first receives the θ_h shock for two periods, and then the θ_l shock for the remainder of the spell. Note, under the U.S. system, this corresponds to an agent who transitions from employment to unemployment, remains unemployed for two periods, and then transitions to non-participation. Figure 1 examines the consumption profile for this particular string of taste shocks.

First, Figure 1(a) plots the replacement ratio (consumption divided by the wage) implied under the optimal contract, for an agent receiving the aforementioned string of shocks. It decreases initially, while the agent remains unemployed, and once the spell of non-participation begins, the replacement ratio increases before eventually remaining constant. Figure 1(a) also plots the evolution of the optimal replacement rate for progressively larger values of $\underline{w}$. There exist two effects as the value of autarky increases. First, the initial drop in consumption (during unemployment) becomes smaller. Second, the increase in consumption during non-participation becomes larger. Figure 1(b) plots the replacement ratio for only the spell of non-participation, which highlights the second effect.

The preceding examples are for an agent who, under the U.S. system, begins unemployed and transitions to non-participation. The increasing consumption result also applies to an agent who simply starts in non-participation and remains there (i.e. receives θ_l shocks for the entire non-employment spell). Figure 1(c) displays consumption for this agent during the non-employment spell. In this case, consumption is monotonically increasing during the non-employment spell. If the agent starts the non-employment spell in unemployment (i.e. receives the θ_h shock and searches) and remains there, consumption decreases monotonically, which is displayed in Figure 1(d).

It is also important to note that the dynamics of consumption analyzed in this section pertain to an agent receiving a sequence of identical shocks; e.g. a sequence of θ_l shocks. In these cases, the dynamics of consumption are smooth. In general, however, consumption is not smooth, as the agent receives a sequence of varying shocks. Given a sequence of varying shocks, consumption jumps around between relatively high values (θ_h shock) and relatively low values (θ_l shock).

5.6 Quantitative Relevance of Theoretical Results

Proposition 1 implies increasing consumption for agents who begin a non-employment spell with an initial promised utility in an interval of $\underline{w}$. Thus, the quantitative relevance depends on how many agents actually may expect to have a level of promised utility in this range during a non-employment spell, according to the optimal allocation. Denote by w^* the level of promised utility such that increasing consumption occurs for $w_0 \in [\underline{w}, w^*)$. According to the equilibrium distribution of w, 76% of the non-employed have a level of expected utility in this critical range where the results of Proposition 1 apply. As the value of autarky increases (i.e. $\underline{w}$), the fraction of the population in the relevant range increases. Table 3 compares this fraction for different levels of autarky. Furthermore, Figure 2 plots the steady state distribution over w for several different values of $\underline{w}$. As $\underline{w}$ increases, the distribution becomes more concentrated at the lower bound, increasing the fraction in the relevant range to nearly 100%.

What does the theoretical result tell us about the world we live in? First, it is reasonable to assume that there are participation constraints (i.e. $\underline{w}$) for the UI authority. That is, they may only "punish" an agent up to some threshold, beyond that level, the agent prefers autarky. Agents never have to participate in the program.[12] Second, to effectively provide incentives to exert effort (searching or retaining the current job), the UI authority needs to elicit an individual's desire to exert effort; i.e. their taste shock. In order to provide the proper incentives to truthfully reveal their desire to remain in the labor force, while simultaneously satisfying the participation constraint, for some agents consumption must increase over the duration of non-employment. This policy differs significantly from the current U.S. system (which suggests constant consumption) and from a standard moral hazard model (that suggests consumption should decrease monotonically).

6 Welfare Analysis

To determine the quantitative significance of adverse selection, consider a comparison of welfare under the current U.S. system relative to an allocation from a planner who ignores adverse selection and focuses only on moral hazard. After calculating the welfare gains of switching from the current U.S. system to the moral hazard only allocation, one can determine how much more is gained by considering both adverse selection and moral hazard.

6.1 Naïve Planner

The comparison is the following: consider a planner, referred to as the *naïve planner*, who provides an allocation that only attempts to satisfy incentive compatibility associated with moral

[12]Indeed, data on the U.S. unemployment insurance system suggests only a fraction of those eligible for UI actually participate. For example, see Blank and Card (1991) and Auray et al. (2013).

hazard, so the allocation is independent of the shock an agent receives. Specifically, the naïve planner does not recognize that agents receive a taste shock $\theta \in \{\theta_l, \theta_h\}$, but instead believes preferences are based on $E(\theta)$.

The naïve planner offers a contract to agents, specifying current period utility, $\tilde{v}_j(i)$, and future promised utility, $\tilde{w}_j(i)$, for $j = e, u$ and $i = e, u$.[13] The cost of providing this contract, to an agent in employment state j last period, is denoted by $\tilde{G}_j(w)$. Since the naïve planner believes preferences are based on $E(\theta)$, the promise keeping and incentive constraint are given by (see the Online Appendix for details of the naïve allocation):

$$w = \pi_j E(\theta)\left[\tilde{v}_j(e) + (1-\pi_j)\tilde{v}_j(u)\right] - \nu + \beta\left[\pi_j \tilde{w}_j(e) + (1-\pi_j)\tilde{w}_j(u)\right] \tag{22}$$

$$\pi_j E(\theta)\left[\tilde{v}_j(e) + (1-\pi_j)\tilde{v}_j(u)\right] - \nu + \beta\left[\pi_j \tilde{w}_j(e) + (1-\pi_j)\tilde{w}_j(u)\right] \geq E(\theta)\tilde{v}_j(u) + \beta\tilde{w}_j(u) \tag{23}$$

The naïve allocation represents a standard repeated moral hazard model of unemployment insurance, and a well known result in these models is the incentive constraint for effort binds at the optimal solution. Thus, when faced with the naïve planner's allocation, the agent receiving the θ_h shock strictly prefers exerting effort, while the agent receiving the θ_l shock prefers shirking. As a result, the cost of providing this allocation, denoted G_j^{naive}, is given by:

$$G_j^{naive}(w) = \mu\tilde{G}_j(w) + (1-\mu)\left[C(\tilde{v}_j(u)) + \beta\tilde{G}_u(\tilde{w}_j(u))\right] \tag{24}$$

It is important to note, despite costs not following $\tilde{G}_j(w)$, unobservable effort implies the naïve planner cannot determine that an agent receiving the shock θ_l shirked.

[13]The naïve planner is assumed to prefer that an agent searches while unemployed. As in the case of the optimal contract in Section 4, for higher levels of promised utility, the naïve planner prefers not to have agents exert search effort. Again, the fully general problem in this case requires lotteries over effort choices. The most general version is omitted here, as this particular allocation is used for quantitative purposes, and the more general case remains irrelevant for the parameters considered here.

6.2 Welfare Comparison

Now consider the welfare delivered by the current U.S. unemployment insurance system. Given the implied decision rules for effort from (15)-(16), the expected welfare under the current U.S. system can be calculated. Recall that under the U.S. system, taxes are determined to balance the budget. Denote the expected lifetime utility delivered under this system as w_{US}. Thus, the relevant welfare comparison is to calculate the level of w such that the optimal contract has a zero expected present value cost.[14] Denote this level of utility as w_P (w_N for the naïve planner).[15] Then, the welfare comparison is made in consumption equivalent terms according to (where W denotes the welfare gain/loss):

$$W = \exp\left[\left(\left(w_P - w_{US}\right)(1-\beta)\right] - 1 \tag{25}$$

Table 4 displays these welfare comparisons under both the optimal contract and the naïve planner's allocation.

For the baseline parametrization, switching from the current U.S. system to the naïve planning allocation increases welfare by 1.06% in consumption equivalent terms. While the naïve planner's allocation provides a significant welfare gain relative to the U.S. system, the optimal contract achieves a welfare gain from the U.S. system of 1.94%; an additional 0.87% over the naïve planner, as the fourth column in Table 4 shows. The additional gains achieved by the optimal contract potentially arise from two sources: an "output" effect and a "consumption" effect. The output effect

[14] An alternative is to use the average level of expected utility implied by the equilibrium distribution of w. The issue with this comparison is that the naïve planner's equilibrium generally involves a different q and thus a different distribution over w. Thus, the comparison between the two allocations involves distributional issues. Fixing q keeps the comparison strictly over differences in consumption allocations.

[15] Actually, the naïve planner's allocation delivers expected utility of slightly above w_N. This occurs because an agent receiving θ_l does not exert effort, and as a result receives a higher level of expected utility from the allocation than the planner thought he was delivering.

occurs because the optimal contract may induce effort from all agents, while the naïve planner's allocation only induces effort from an agent receiving θ_h. The second source of welfare gains, the "consumption" effect, arises because the optimal contract provides a more efficient allocation of consumption. The last two columns of Table 4 present the magnitude of these effects, respectively.

6.3 Heterogeneity

Transitions to non-participation represent an important aspect of the quantitative analysis. The baseline cases presented above pertain to transitions calculated for the entire population. Of course, labor market attachment, and thus the transitions between labor market states, may vary significantly for different segments of the population. This section presents the same welfare analysis as above, but for transition rates calculated on different sub-samples of the population.

Specifically, consider the transitions for different age groups: 16-24, 25-54, and 55 and older. Note, these represent observable characteristics that the planner can condition allocations on. To gauge each groups' "attachment" to the labor force, only the fraction of non-participants in each group that want a job now are considered. From Section 5.3, recall this fraction is 7.2% for the general population. For the youngest cohort, attachment is greatest, with 12.3% of non-participants wanting a job now. This is compared to 11.2% for 25-54, and 2% for the 55 and older group. The probability $1 - \mu$ multiplies this percent by the transition probability from unemployment to non-participation.

Table 5 presents the welfare gains for each group, as well as the differences in transition rates between labor market states.[16] There exist large additional gains from the optimal contract for the youngest cohort. This group has both the highest benefit to searching (i.e. highest value of π_u)

[16] The comparison for each group is made by finding the lowest value of autarky ($\underline{w}$) that delivers a balanced budget equilibrium (i.e. q close to β). This allows a consistent comparison for all groups, and with the baseline case.

and the most "attachment" to the labor market (highest percentage of non-participants who want a job now). The lowest additional gains occur for the oldest cohort (55 and older), who have the lowest benefit to searching and the least attachment to the labor force. Indeed, in the planner's equilibrium for the oldest cohort, those agents with higher levels of promised lifetime utility are not required to search if non-employed; i.e. the transition from unemployment to non-participation is efficient for some agents in this group.

6.4 Robustness

There exist many potential alternative specifications to the planning problem studied in this paper. This section examines whether the baseline quantitative results are robust to these potential alternatives. Three alternatives are considered. The Online Appendix presents the details of these alternative specifications. This section summarizes the quantitative results. Table 6 displays the welfare gains in each formulation, including the baseline case.

First consider the alternative specification of preferences discussed in Section 3.3. That is, per-period utility is given by $\mathrm{v}(c_t) - \theta_t \nu$. Since the consumption effect in Table 4 is relatively large, one may be concerned that the welfare gains are artificially inflated by the particular preference specification. The second row of Table 6 displays the welfare gains for the alternative preferences. Relative to the baseline case (first row), the additional welfare gains are essentially unchanged, and actually increase slightly. This suggests that quantitatively, allowing the taste shock to directly impact the marginal utility of consumption does not significantly affect the quantitative results.

Second, some of the additional welfare gains achieved by the optimal contract arise due to the taste shocks affecting employed agents. Since these may not be appropriate to include in an analysis of unemployment insurance, the following robustness check is provided. Consider a version of the planning problem similar to that in Hopenhayn and Nicolini (1997), where once employed agents no longer receive taste shocks. They incur a utility cost to working, ν, but do not have a choice whether to exert the effort or not; with exogenous probability $1-\pi_e$ they lose the job and transition to non-employment. Preferences and all non-employed states are as in the baseline case.

The third row in Table 6 displays the welfare results in this case, and the fourth row refers to this "Hopenhayn and Nicolini" set-up under the alternative specification of preferences. Surprisingly, the welfare gains (both total and extra) are larger in the Hopenhayn and Nicolini case. This owes to the fact there is no incentive problem for employed agents, so the planner does not incur the costs of providing incentives for these agents.

7 Conclusion and Discussion

The results in this paper suggest strong implications for current UI policy. First, "unemployment" insurance should be extended to those who are currently non-participants (but still want a job and are available) under the U.S. system. Moreover, according to the optimal contract, these non-participants should be searching (i.e. do not transition to non-participation). Second, in terms of implementation, consider the following. Each time a worker files a claim for benefits, they should report their desire to remain in the labor force or not (i.e. search or not search). Depending on this report, the UI authority offers a benefit scheme. In equilibrium, the optimal allocation implies that repeated reports of a desire to exit the labor force imply an increasing consumption profile. The agent should report this desire each period, and the allocation depends on the entire history of reports and the entire history of employment. As others have noted (e.g. Wang and Williamson (2002)) this history dependence could be implemented via a scheme of UI accounts; the balance in the account increases and decreases according to actual histories. This device serves to act as the promised lifetime utility variable in the optimal contract. The quantitative results show that there exist large welfare gains from considering the report of an agent's desire to continue searching.

There are several aspects of my analysis that offer interesting directions for future research. First, the quantitative analysis primarily illustrates how the inclusion of dynamic adverse selection changes the qualitative predictions of the dynamics of benefits in a model of optimal UI. While it uses a subset of the unemployment to non-participation group to quantify the effects of moral hazard, this clearly does not take into account the full heterogeneity of the population. Future research should indeed focus on incorporating more persistent shocks that would allow such additional heterogeneity to be analyzed.

Second, the policy prescriptions for non-participants rests on being able to identify the 7.2% of non-participants who want a job and are able and available for one. That is, while the taste shock remains unobservable, the analysis assumes that the UI agency can offer a policy that specifically targets this subset of non-participants. In many cases this assumption is relatively innocuous. Those who do not want a job now, or are unable and unavailable for one, are often easily identifiable. Examples include those disabled, those engaged in child care, or full-time students. In such cases, the UI agency can easily identify these types as not in the able and available group.

It should be noted, however, that the policy prescriptions presented in this paper do come with this caveat. Indeed there may exist some portion of non-participants that are not easily identifiable as in or out of the able and available group. This also underscores the need for future research to incorporate more persistent shocks, which would allow the policy prescriptions to be less dependent on identifying this subset of non-participants.

References

Andolfatto, D., Gomme, P., 1996. Unemployment insurance and labor-market activity in Canada. Carnegie-Rochester Conference Series on Public Policy 44, 47–82.

Atkeson, A., Lucas, R., 1992. On efficient distribution with private information. Review of Economic Studies 59, 427–453.

Atkeson, A., Lucas, R., 1995. Efficiency and inequality in a simple model of efficient unemployment insurance. Journal of Economic Theory 66, 64–88.

Auray, S., Fuller, D. L., Lkhagvasuren, D., 2013. Unemployment insurance take-up rates in an equilibrium search model. Working paper, Concordia University.

Blank, R., Card, D., 1991. Recent trends in insured and unisured unemployment: Is there an explanation? The Quarterly Journal of Economics 106, 1157–1189.

Davidson, C., Woodbury, S., 1998. The optimal dole with risk aversion and job destruction. Working paper, Michigan State University.

Engelhardt, B., Fuller, D. L., 2012. Efficient labor force participation with search and bargaining. Labour Economics 19, 388–402.

Farhi, E., Werning, I., 2007. Inequality and social discounting. Journal of Political Economy 115, 365–402.

Garibaldi, P., Wasmer, E., 2005. Equilibrium Search Unemployment, Endogenous Participation, and Labor Market Flows. Journal of the European Economic Association 3, 851–882.

Gautier, P. A., Moraga-Gonzalez, J. L., Wolthoff, R. P., 2009. Structural estimation of search intensity: Do non-employed workers search enough? Discussion paper no. 07-071/3, Tinbergen Institute.

Hagedorn, M., Kaul, A., Mennel, T., 2010. An adverse selection model of optimal unemployment insurance. Journal of Economic Dynamics and Control 34, 490–502.

Hall, R. E., Milgrom, P. R., 2008. The limited influence of unemployment on the wage bargain. American Economic Review 98, 1653–1674.

Hopenhayn, H., Nicolini, J., 1997. Optimal unemployment insurance. Journal of Political Economy 105, 412–438.

Hopenhayn, H., Nicolini, J., 2009. Optimal unemployment insurance and employment history. Review of Economic Studies 76, 1049–1070.

Kocherlakota, N., 2004. Figuring out the impact of hidden savings on optimal unemployment insurance. Review of Economic Dynamics 7, 541–554.

Meyer, B., 1990. Unemployment insurance and unemployment spells. Econometrica 58, 757–782.

Mitchell, M., Zhang, Y., 2010. Optimal unemployment insurance with hidden trade. Journal of Economic Theory 145, 2078–2107.

Pavoni, N., 2007. On optimal unemployment compensation. Journal of Monetary Economics 54, 1612–1630.

Pavoni, N., 2009. Optimal unemployment insurance, with human capital decreciation,and duration dependence. International Economic Review 50, 323–362.

Pries, M., Rogerson, R., 2009. Search frictions and labor market participation. European Economic Review 53, 568–587.

Shavell, S., Weiss, L., 1979. The optimal payment of unemployment insurance benefits over time. Journal of Political Economy 87, 1347–1362.

Shimer, R., 2012. Reassessing the ins and outs of unemployment. Review of Economic Dynamics 15, 127–148.

Shimer, R., Werning, I., 2006. On the optimal timing of benefits with heterogeneous workers and human capital depreciation. Working paper 06-12, MIT.

Shimer, R., Werning, I., 2008. Liquidity and insurance for the unemployed. American Economic Review 98, 1922–1942.

Thomas, J., Worral, T., 1990. Income fluctuations and assymetric information: an example of repeated principal-agent problem. Journal of Economic Theory 51, 367–390.

Wang, C., Williamson, S., 2002. Moral hazard, optimal unemployment insurance, and experience rating. Journal of Monetary Economics 49, 1337–1371.

Werning, I., 2002. Optimal unemployment insurance with hidden savings. Working paper, University of Chicago.

Table 1: Reason for not searching (as a % of group)

Discouragement over job prospects	28%
Reasons other than discouragement	72%
Family responsibilities	10%
In school or training	14%
Ill health or disability	8%
Other (includes child care and transportation)	40%

Source: www.bls.gov. Data refers to the subset of non-participants who want a job now, and are able and available to work. The table displays the reason why these individuals entered non-participation, and the percentage of this group in each category. The responses are broken into two main groups: those discouraged, and those not. For those not searching for reasons other than discouragement, the table gives several possible responses.

Table 2: Transition rate comparisons

Model	U-E	U-O
U.S. System	0.32	0.021
Naïve Planner	0.32	0.021
Optimal Contract	0.33	0

Notes: Transition rates are quarterly averages of monthly data. They represent the probability of transitioning from unemployment to employment (U-E) and from unemployment to non-participation (U-O). The U-O transition (in the data) reflects the rate for only those who want a job now and are available for one. Since all non-employed search in the optimal allocation, the U-O transition probability is zero.

Table 3: Fraction of non-employed with $w \in [\underline{w}, w^*)$

d	$\underline{w}$	w^*	Fraction
Baseline	34%	36%	0.758
0.26	35%	37%	0.884
0.30	39%	41%	0.996
0.35	43%	45%	0.998

Notes: This table describes the fraction of non-employed agents who have a level of expected utility (in the steady state distribution) in the range where the increasing consumption result applies. The first column displays the value of autarky consumption (as a % of the wage). The second displays the lower bound on future utility promises, in per-period consumption as a percentage of the wage: $\exp[(1-\beta)\underline{w}]/y$. The third column reports the value of w where the policy function $w_u^l(u)$ crosses the 45 degree line. The last column gives the fraction of non-employed with a promised utility, w, less than this value, in the steady state distribution.

Table 4: Welfare gains (in %), relative to current U.S. system

d	$\frac{\theta_h}{\theta_l}$	Total Gain	Naïve	Extra Gain	Output	Consumption
Baseline	3.49	1.94	1.06	0.87	0.06	0.81
0.30	3.03	1.71	0.72	0.99	0.06	0.93
0.35	2.64	1.16	0.39	0.77	0.06	0.71
0.37	2.50	1.07	0.15	0.91	0.06	0.85

Notes: All welfare gains are consumption equivalent gains. The first column displays different levels of the lower bound, via d, autarky consumption. The second column displays the ratio of taste shocks. The third column gives the welfare gain achieved by the optimal contract, and the fourth column the gain from the naïve allocation, both relative to the current U.S. system. The last three columns detail the additional gains from the optimal contract, and their source, respectively. "Extra" refers to the total additional gains achieved relative to the naïve planner, "Output" and "Consumption" detail the magnitudes of each effect, respectively.

Table 5: Optimal Policy Comparison Across Age Groups

Group	π_e	π_u	$1-\mu$	$\frac{\theta_h}{\theta_l}$	ν	Extra Gain	Total Gain
Baseline	0.982	0.33	0.021	3.49	0.98	0.87	1.94
16-24	0.96	0.36	0.049	3.19	0.78	1.87	3.23
25-54	0.986	0.32	0.028	3.17	0.92	1.10	1.77
55 and older	0.99	0.29	0.007	3.33	1.00	0.30	0.99

Notes: The first column details the age group considered. The next five columns summarize the changes in parameters for each group. For each age group the table displays the relevant parameters: probability of retaining employment (π_e), job finding probability (π_u), unemployment to non-participation (want a job and available) transition probability $(1-\mu)$, the corresponding ratio of taste shocks required ($\frac{\theta_h}{\theta_l}$), and the utility cost of effort (ν). The last two columns report the additional gains achieved by including adverse selection (relative to the naïve planner) and the total gains achieved by the optimal contract (relative to the U.S. system), respectively. The additional gains achieved by the optimal contract decrease with the age considered, as attachment to the labor force decreases.

Table 6: Robustness: Welfare gains (in %) under alternative specifications

Specification	Total Gain	Naïve	Extra Gain
Baseline	1.94	1.06	0.87
Alternative Preferences	1.59	0.55	1.04
Hopenhayn-Nicolini	2.45	0.52	1.92
H-N Alt. Preferences	1.35	0.55	0.79

Notes: The table displays the consumption equivalent welfare gains for alternative specifications of the planning problem. The first column describes the model version and the second gives the gains achieved by the optimal contract (relative to the current U.S. system). Finally, the third and fourth columns display the gains under the naïve allocation and the additional gains from considering adverse selection, respectively. All specifications have welfare gains of similar magnitudes, with the highest extra gain coming in the Hopenhayn and Nicolini specification.

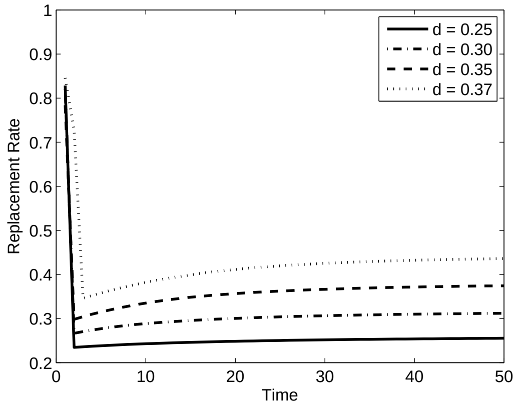

(a) Optimal contract replacement ratios

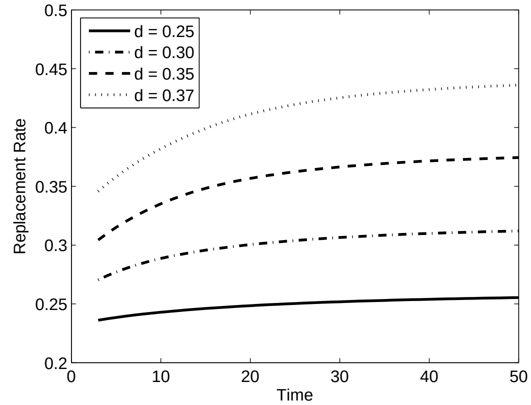

(b) Optimal contract, non-participation portion

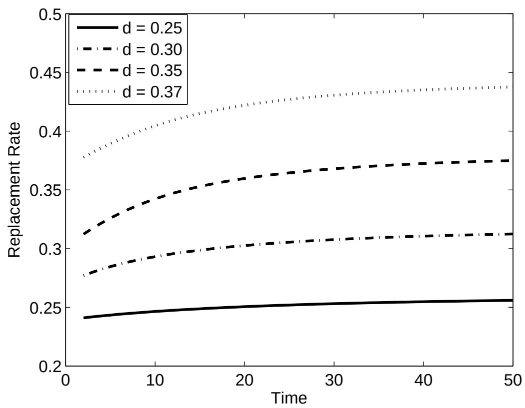

(c) Optimal contract, string of only θ_l

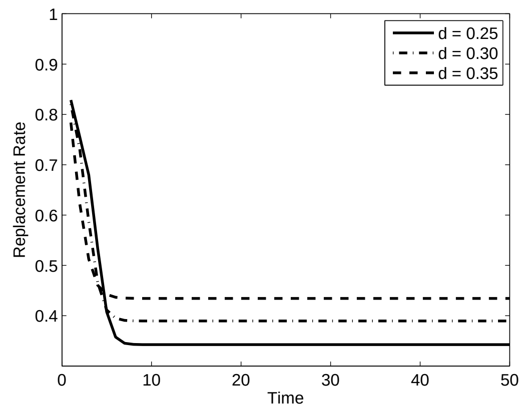

(d) Optimal contract, string of only θ_h

Figure 1: Consumption profile, spell of non-employment. Each graph plots the dynamics of UI benefits for different portions of the non-employment spell and combinations of the taste shock. The vertical axis plots the "replacement rate," which is consumption divided by the wage. The upper left figure displays the dynamics of the replacement rate for a sequence of two periods of unemployment (θ_h shock) followed by non-participation (repeated θ_l shocks). The upper right figure displays only the non-participation portion of this spell. The lower left and right figures plot the replacement rate over a spell of *only* non-participation (θ_l shocks) and unemployment (θ_h shocks), respectively. In all experiments, the agent begins with the level of promised utility delivered to an employed agent by the current U.S. system.

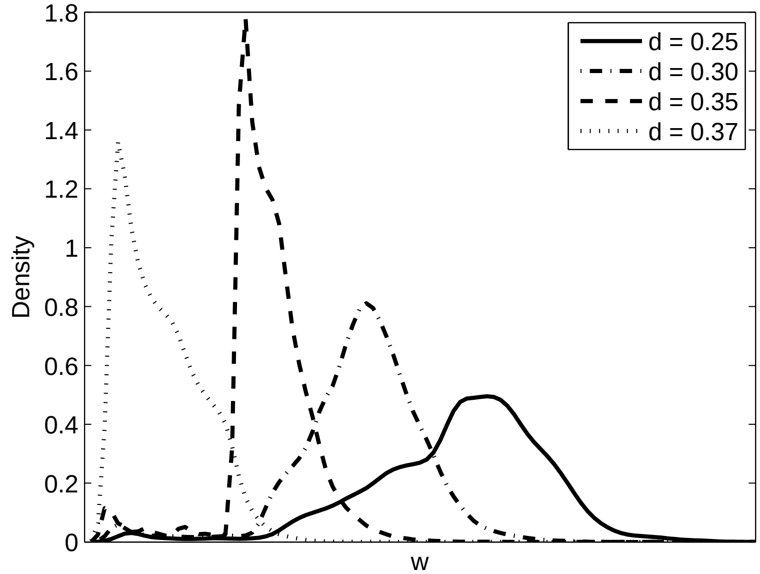

Figure 2: Expected utility distribution, varying $\underline{w}$. Each curve represents the steady state distribution over w that obtains in the planning problem. The support of each distribution varies, they are overlayed for visual comparison. As $\underline{w}$ increases (i.e. d increases), the distribution becomes more concentrated and shifts to the left.